

Old Format

The CDF-1 file format is the default file format which has been in use through netCDF library version 3.5.1. The definitive resource for these limits can be found at Unidata's [NetCDF Classic Format Limitations](#) page

Extending Limits

The CDF-2 file format relaxes many of the above limitations. In order to create a CDF-2 formatted file, you must pass the flag NC_64BIT_OFFSET when creating the dataset. Parallel netCDF supports the CDF-2 file format as well, though shares the same limits as serial NetCDF. Read more at UniData's [64-bit Offset Format Limitations](#) page for limitations on the number of variables allowable per file, maximum sizes of record and non-record variables, and exceptions. However, in CDF-2, arrays are still limited to a maximum of 2^{32} elements.

The CDF-5 file format further relaxes the limitation to allow arrays with more than 2^{32} elements. To create a CDF-5 file, the mode NC_64BIT_DATA must be used when creating the file. Please note that at the time of this writing (Oct.15, 2012), UniData's NetCDF has not supported CDF-5 yet.

Record vs. Non-Record

One subtle point about the CDF file format limitations is that record variables (those with one NC_UNLIMITED dimension) can be quite large. The product of the fixed-sized dimensions of a record variable defines the size of a single record. There can be quite a few of these records in a record variable, though, so that can be one way around the file format limitations on record size.

The drawback of using record variables is that with multiple record variables, the data for each variable are interleaved on a per-record basis. This data layout makes it much harder to figure out the optimal access pattern. Ideally, though, programs would read and write record variables one record at a time.